

ACCURACY OF THE ERBS DEFINITIVE ATTITUDE DETERMINATION SYSTEM IN THE PRESENCE OF PROPAGATION NOISE*

D. Chu and E. Harvie
Computer Sciences Corporation

ABSTRACT

Definitive attitude solutions are supposed to be the most accurate possible. For the Earth Radiation Budget Satellite (ERBS), this has been accomplished by using gyro rates to transform many nonsimultaneous observations to a common time point and then averaging to reduce the effects of observation noise. Rate quality is critical to realizing improved accuracy with this method. Gyro deterioration, which shows up as large observation residuals and discontinuities between contiguous batch solutions, now discourages using the batch approach for ERBS. To address this problem, a simple Kalman filter is tried in place of the batch estimator. The filter works well as long as the attitude is completely observable. During periods without Sun coverage, however, the extrapolated yaw may diverge and then change abruptly when the Sun returns to the sensor field of view. Causes of this behavior are discussed, and some solutions are tried that address the observability aspect of the problem.

*This work was supported by the National Aeronautics and Space Administration (NASA)/Goddard Space Flight Center (GSFC), Greenbelt, Maryland, Contract NAS 5-31500.

1. INTRODUCTION

Definitive attitude solutions are to be the most accurate possible and are to be continuous over the entire mission. Under the best conditions, this can be difficult both technically and operationally (Reference 1). In this paper, we consider one of the technical problems—that caused by attitude propagation error.

As a spacecraft gets old, its gyros may become noisy. This has been the case for the Earth Radiation Budget Satellite (ERBS). Without accurate angular velocity measurements, its definitive attitude determination system has been unable to provide continuous solutions or fit all observations.

To remedy this, the batch estimator is replaced by a Kalman filter that accounts for the gyro noise. The results obtained are described, and recommendations are made for further work to improve the approach.

2. ERBS BACKGROUND

ERBS is an Earth-pointing spacecraft in a 57-degree (deg) inclination, nearly circular orbit with a period of 97 minutes. It has two horizon sensors for measuring geodetic pitch and roll plus two Sun sensors whose alpha angles measure yaw. In addition, ERBS has a three-axis magnetometer. The horizon sensors are accurate to 0.5 deg, due primarily to horizon radiance effects. The Sun sensor is accurate to 0.05 deg, and the magnetometer is accurate to 3 deg. Because of its lower accuracy, the magnetometer is not used for definitive attitude determination. ERBS also has two redundant three-axis gyros for measuring its angular velocity. Both gyros have now partially failed. The first failure occurred on a pitch-axis gyro in August 1986, 22 months after launch. The second gyro failed on a pitch axis as well in July 1988. At present, only roll and yaw rates are available.

While the horizon sensors always see the Earth, the Sun sensors see the Sun for only part of each orbit. This means that attitude is not completely observable at all times, and good angular velocity measurements are needed to provide a complete history of the attitude. This incomplete observability also necessitates special care in dividing up each day's data for computing attitude.

3. DEFINITIVE ATTITUDE DETERMINATION

The ERBS definitive attitude system has three parts that determine the effect of propagation error—the estimator, the segmenter, and the smoother (Reference 2). The current batch least squares estimator updates its previous estimate of the state (epoch attitude, gyro bias, and scale factor errors) using all the observations at once. Sensor observations from different times are in effect propagated back to one epoch time, where they are averaged to reduce the effect of measurement noise.

If propagation were perfect, any number of observations could be so transformed, giving a solution of unlimited accuracy. In practice, however, propagation does add uncertainty,

which increases with the time between the epoch and the observation times. Thus, the residuals, or differences between the observed and predicted values, reflect both sensor noise, which is stationary, and propagation error, whose uncertainty grows with time.

If these residuals grow linearly with time from epoch or in proportion to the gyro output rates, they can be reduced by solving for rate bias and scale factor corrections along with the epoch attitude. Since bias errors grow with time, knowing the attitude at both ends of the batch affords the best opportunity for estimating biases. For this reason, the segmenter divides the day's data so that there is Sun sensor coverage at the start and end of every batch. This ensures that the rate biases can be accurately determined.

The smoother also serves to reduce the evidence, if not the effect, of propagation error. Because not all gyro errors are eliminated by bias and scale factor corrections, the time dependence of the attitude history may still be incorrect. No matter what the choice of epoch attitude, the solution cannot be correct over the entire batch. This shows up in the definitive attitude as discontinuities between batches. The smoother discards points around the junction and replaces them with values obtained by linearly interpolating between the new end points as in Figure 1. This is done without regard to sensor observations, and may not significantly improve the solution.

The batch estimator, by ignoring this random part of propagation error, creates a problem that the segmenter cannot fix and the smoother simply covers up.

4. MODELING PROPAGATION ERROR

Although propagation error can cause batch junction discontinuities, it remains to be shown that the ERBS gyro deterioration actually produces such discontinuities. To make this connection, a simplified Farrenkopf gyro noise model is used to predict discontinuity size based on the observations of gyro noise. These predictions are then shown to track the daily average pitch, roll, and yaw junction discontinuities.

Unlike the white noise seen in the sensor observations, errors in the propagated attitude tend to change continuously. For this reason, the propagated angle (θ) and its error ($\Delta\theta$) are modeled as a random process. The model in Figure 2 is a linear second-order system with four zero mean value sources of error:

- Float torque derivative noise (n_u)
- Float torque noise (n_v)
- Initial attitude error ($\Delta\theta(0)$)
- Drift rate bias (b)

The float torque derivative noise is integrated to give a random walk that is added to the true angular rate (ω). A constant, but imprecisely known bias and the float torque noise are also added at that point. At the output, another constant error is added representing the initial or epoch attitude error.




Figure 1. Smoothed Junction Discontinuity




Figure 2. Gyro Noise Model

Because the four error sources have zero mean value, the expected propagation error is zero mean as well. The variance of the error, however, grows with propagation time (t) according to the expression

$$\sigma_{\Delta\theta}^2(t) = \sigma_{\Delta\theta}^2(0) + \sigma_v^2 t + \sigma_b^2 t^2 + \sigma_u^2 t^3/3 \quad (1)$$

The individual error source variances are defined as follows where E is the expectation operator and δ is the Dirac delta function.

$$E[b^2] = \sigma_b^2 \quad (2)$$

$$E[\Delta\theta^2(0)] = \sigma_{\Delta\theta}^2(0) \quad (3)$$

$$E[n_v(t) n_v(\tau)] = \sigma_v^2 \delta(t - \tau) \quad (4)$$

$$E[n_u(t) n_u(\tau)] = \sigma_u^2 \delta(t - \tau) \quad (5)$$

Standard deviations of the ERBS gyro rates were compiled over the course of the mission and provide an estimate of the gyro noise. Figure 3 shows this history of increasing noise.

These standard deviations for the sampled angular velocity ($\sigma_{\dot{\Delta\theta}}$) give an estimate of the float torque noise standard deviation (σ_v) that can be used to predict a typical batch junction discontinuity. Because the sampled rates are actually averaged over the sampling interval (T), their variance is smaller than the float torque noise variance by a factor of 1/T.

$$\sigma_{\dot{\Delta\theta}}^2 = \sigma_v^2/T \quad (6)$$

Since T is 1 second, the numerical values of $\sigma_{\dot{\Delta\theta}}$ and σ_v are equal, but their units differ as they should.

5. EVIDENCE OF GYRO NOISE

Figure 4 plots daily average junction discontinuities and observation residuals against corresponding gyro rate standard deviations. In all cases, the discontinuities and residuals increase with increasing noise.

One thing that requires explanation here is the very steep slope of the yaw and Sun A curves. This is due to the low noise on the yaw gyro and the coupling of propagation




Figure 3. Gyro Rate Standard Deviations

errors on the roll and yaw axes. The large roll gyro noise causes yaw as well as roll propagation errors. Plotting yaw error against yaw gyro noise alone is therefore misleading. A better way to show the dependence might be to combine the roll and yaw into a single curve.

The predicted values of junction discontinuity are based on Equation 1 and use the gyro rate standard deviation for σ_v . Although inaccurate, the predictions are of the same order of magnitude as the observed discontinuities. Also, the rate standard deviation is an overestimate of the gyro noise since it includes the actual angular acceleration of the spacecraft. With a smaller assumed value for σ_v , the predictions would be closer to the observed discontinuities.





Figure 4. Discontinuities and Residuals

6. THE SIMPLE KALMAN FILTER

The fact that pitch should be so hard to estimate illustrates one weakness of the batch estimator. The reason it is so susceptible to propagation error is that it accepts the angular rates almost at face value. After correcting for biases and scale factors, the propagation is assumed to be perfect and observations far from the epoch are given equal weight with those close to the epoch.

The Kalman filter differs from the batch estimator in that it estimates the attitude at each time point rather than at an epoch time (Reference 7). It also gives greater weight to observations made close to the solution time. This seemed to be the answer to the propagation error/discontinuity problem. Old observations would be forgotten, giving the solution more freedom to follow recent observations. Even more, solutions could be made continuous at batch boundaries by starting with the previous batch solution and covariance.

Although different in approach from the batch estimator, the filter makes many of the same computations. For a very simple filter estimating attitude and constant bias, only three major modifications are required to the ERBS batch estimator.

1. Change the partial derivatives

- a. In the batch estimator, the partial derivatives of the current observation residual vector Δy_j

$$\Delta y_j = y_j (\text{observed}) - y_j (\text{computed}) \quad (7)$$

with respect to the current attitude error vector \underline{a}_j are postmultiplied by the epoch-to-current-time propagation matrix $\phi(t_j, t_0)$. This multiplication is omitted in the filter

$$\frac{\partial y_j}{\partial \underline{a}_0} = \frac{\partial y_j}{\partial \underline{a}_j} \phi(t_j, t_0) \rightarrow \frac{\partial y_j}{\partial \underline{a}_j} \quad (8)$$

- b. Filtered gyro biases \underline{b}_j act only over individual time steps before being updated with the bias error vector \underline{a}_j . Rather than accumulate derivatives of the current attitude error with respect to gyro biases from epoch to the current time, as in the batch estimator, the filter computes them only for the most recent time step.

$$\frac{\partial \underline{a}_j}{\partial \underline{a}_j} = \psi(t_j, t_{j-1}) + \phi(t_j, t_{j-1}) \frac{\partial \underline{a}_{j-1}}{\partial \underline{a}_{j-1}} \rightarrow \psi(t_j, t_{j-1}) \quad (9)$$

The matrix $\psi(t_j, t_{j-1})$ comes from integrating the propagation matrix and assumes the bias is constant over this interval.

$$\psi(t_j, t_{j-1}) = \int_{t_{j-1}}^{t_j} \phi(t_j, \tau) d\tau \quad (10)$$

- c. To simplify notation, the combination of the attitude and gyro bias error vectors are referred to as the state \underline{x}_j

$$\underline{x}_j^T = [\underline{a}_j \ : \ \dot{\underline{a}}_j] \quad (11)$$

and the partial derivatives of the observation residuals with respect to the state is called F_j

$$F_j = \left(\frac{\partial \underline{y}}{\partial \underline{a}} \ : \ \frac{\partial \underline{y}}{\partial \dot{\underline{a}}} \right) = \left(\frac{\partial \underline{y}}{\partial \underline{x}} \right)_j \quad (12)$$

2. Update the state at each time step—The batch estimator weights, transforms, and accumulates the observation residuals over the entire batch to give

$$FTWY \equiv \sum_{i=1}^N F_i^T W \Delta y_i \quad (13)$$

The observation weight matrix W is diagonal with elements equal to the reciprocal of the observation variances

$$W = \begin{pmatrix} 1/\sigma_{y1}^2 & 0 \\ 0 & 1/\sigma_{ym}^2 \end{pmatrix} \quad (14)$$

The batch estimator also accumulates the normal matrix, which is the inverse of the covariance matrix P .

$$\sum_{i=1}^N F_i^T W F_i + W_0 \equiv P_N^{-1} \quad (15)$$

The a priori estimate weight matrix W_0 is also diagonal and has as its elements the reciprocal of the a priori state variable variances

$$W_0 = \begin{pmatrix} 1/\sigma_{y1}^2 & 0 \\ 0 & 1/\sigma_{yn}^2 \end{pmatrix} \quad (16)$$

In batch estimation, the epoch state estimate is updated after all the observation residuals have been accumulated.

$$\Delta \underline{x}_0 = (FTWF)^{-1} (FTWY) = (P_N) (FTWY) \quad (17)$$

The filter makes this correction at each time step using the weighted transformed observation residuals for that instant and the accumulated covariance matrix.

$$\Delta \underline{x}_j = (P_j) (F_j^T W \Delta y_j) \quad (18)$$

3. Propagate the covariance matrix—The batch estimator propagates the attitude and predicts observations at each time step. The filter propagates the covariance matrix as well. This is done by pre- and postmultiplying the covariance matrix by the incremental state propagation matrix $\Phi(t_j, t_{j-1})$ and its transpose.

$$P_j = \Phi(t_j, t_{j-1}) P_{j-1} \Phi^T(t_j, t_{j-1}) + Q \quad (19)$$

A contribution Q due to the propagation error is also added at this point (Reference 3).

$$Q = \begin{bmatrix} \sigma_v^2 (t_j - t_{j-1}) + \frac{1}{3} \sigma_u^2 (t_j - t_{j-1})^3 & -\frac{1}{2} \sigma_u^2 (t_j - t_{j-1})^2 \\ -\frac{1}{2} \sigma_u^2 (t_j - t_{j-1})^2 & \sigma_u^2 (t_j - t_{j-1}) \end{bmatrix} \quad (20)$$

The state propagation matrix includes the attitude propagation matrix as its upper left block

$$\Phi(t_j, t_{j-1}) = \begin{pmatrix} \phi(t_j, t_{j-1}) & \psi(t_j, t_{j-1}) \\ 0 & I_{3 \times 3} \end{pmatrix} \quad (21)$$

The attitude parameterization chosen for the filter is the attitude error vector. This representation fits in with the existing batch differential correction and avoids complications from unit length and orthonormality constraints. To update the estimated geocentric inertial-to-body attitude matrix $A_{B/I}$, the error vector is converted to a rotation matrix $R(\underline{a}_j)$ which is used to premultiply the old value of $A_{B/I}$.

$$A_{B/I}(t_j) = R(\underline{a}_j) A_{B/I}(t_{j-1}) \quad (22)$$

The gyro bias corrections \underline{a}_j are simply added to the previous bias values

$$\underline{b}_j = \underline{b}_{j-1} + \underline{a}_j \quad (23)$$

7. COMPARING THE BATCH ESTIMATOR AND THE FILTER

After being implemented, the filter was used with low and high noise gyro rate and was compared to the results from the batch estimator. For timespans having Sun sensor coverage, the attitude is completely observable, and the filter provides smaller root-mean-square (rms) observation residuals. This is shown in Table 1. There is no chance of junction discontinuities because the filter starts with the last attitude of the preceding timespan.

Table 1. RMS Residuals—Complete Observability

Observation Residual Variance (deg²)

Low Noise	Roll	Pitch	Sun A	Overall
Filter	0.16	0.30	0.03	0.87
Batch	0.10	0.31	0.01	0.87
High Noise	Roll	Pitch	Sun A	Overall
Filter	0.05	0.13	0.00	0.17
Batch	0.85	0.73	0.28	1.07

Over timespans without continuous Sun coverage, the filter can diverge and cause jumps even greater than those between batches. An example of such a filter discontinuity is provided in Figure 5. There, the yaw diverges until the Sun comes back into view. The size of the jump depends strongly on the actual gyro noise and the assumed gyro noise, i.e., tuning of the filter.




Figure 5. Drifting of Filtered Yaw

As shown in Table 2, the residuals are still much smaller for the filter than for the batch estimator. This is in spite of the discontinuity. The reason is that the filter fits all the observations better except for the first few following the start of Sun coverage.

Table 2. RMS Residuals—Incomplete Observability

Observation Residual Variance (deg²)

Low Noise	Roll	Pitch	Sun A	Overall
Filter	0.09	0.46	1.26	0.40
Batch	0.31	2.58	0.14	1.50
High Noise	Roll	Pitch	Sun A	Overall
Filter	0.04	0.08	1.45	0.19
Batch	29.66	9.67	0.66	17.37

8. DIVERGENCE REMEDIES

The batch estimator, despite its problems with junction discontinuities and observation residuals, does not drift during periods of incomplete attitude observability. This is because it has access to data from the entire batch at once. Thus, the batch estimator can choose a bias satisfying the observations at both the beginning and the end of the batch. The filter, on the other hand, knows only what has been seen up to the current time and does not have this advantage.

To improve filter performance during periods without Sun, the magnetometer can be used along with the Sun and Earth sensors. Although it is the least accurate sensor, the magnetometer provides some measure of the yaw at all times. As shown in Figure 6, this prevents divergence but does not provide a very accurate definitive attitude solution.

An alternative to improving the observability is to improve the propagation. To do so, the batch estimator can be used to compute a gyro bias over the period lacking Sun coverage. This bias can then be used in the filter to reduce the divergence, as shown in Figure 7.




Figure 6. Yaw Using Magnetometer

A variant on this "two-pass" processing is to average the forward-filtered solution with a backward-filtered solution obtained by starting from the end of the timespan and letting time run backward. This "filter-smoother" can still diverge toward the middle of the Sun coverage gap, but does not change abruptly the way a simple filter can.

The approach that should, perhaps, have been tried first is to "tune" the gyro noise parameter values to reduce the divergence. The values used here for the Q-matrix were based on the noise estimated from the raw gyro rates. It was noted that these assumed values strongly affect the filter divergence, but no effort was made to find values that worked well in all cases. Judging from the satisfactory performance of the ERBS gyro-compass, which uses only the Earth sensors and gyro, it should be possible to find such values. Perhaps the gyro noise levels used here were too large and so caused the filter to forget old observations too quickly.




Figure 7. Yaw With Batch Bias

9. SUMMARY

This paper uses the ERBS definitive attitude determination system to demonstrate two things: first, that gyro noise affects batch attitude accuracy, as seen in the observation residuals and batch junction discontinuities; and second, that Kalman filtering can reduce those residuals and eliminate the junction discontinuities as well. What was also demonstrated, albeit inadvertently, was that during times of incomplete attitude observability, the filtered attitude may diverge. Patchwork remedies were tried, such as using additional sensors or future observations, but the basic cause of the problem seems to have been the "tuning" of the filter. Had a smaller value been assumed for the gyro noise, the attitude would not have diverged as quickly as it did. In spite of the present incomplete success,

Kalman filtering still promises accuracy equal to or better than that of the batch estimator once the tuning process is mastered.

REFERENCES

1. Computer Sciences Corporation, CSC/TR-89/6021, *A Study of Definitive Attitude Determination*, M. Lee and J. Rowe, September 1989
2. --, CSC/SD-82/6013, *Earth Radiation Budget Satellite (ERBS) Attitude Ground Support System (AGSS) Functional Specifications and Requirements*, G. Nair, September 1982
3. R. L. Farrenkopf, "Generalized Results for Precision Attitude Reference Systems Using Gyros," AIAA Paper 74-903, AIAA Mechanics and Control of Flight Conference, Anaheim, CA, August 1974
4. J. R. Wertz, *Spacecraft Attitude Determination and Control*, D. Reidel Publishing Co., 1978, p. 269
5. O. Filla, T. Z. Willard, and D. Chu, *In-Flight Estimation of Gyro Noise*, Flight Mechanics/Estimation Theory Symposium, NASA-Goddard Space Flight Center, May 1990
6. E. Harvie, D. Chu, and M. Woodard, "The Accuracy of Dynamic Propagation," Flight Mechanics/Estimation Theory Symposium, NASA-Goddard Space Flight Center, May 1990
7. E. J. Lefferts, F. L. Markley, and M. D. Shuster, "Kalman Filtering for Spacecraft Attitude Estimation," *Journal of Guidance*, September-October 1982, vol. 5, no. 5, pp. 417-429